

Curiosity-Aware Bargaining

Cédric L.R. BURON ^{*1,2}, Sylvain DUCTOR², and Zahia GUESSOUM^{2,3}

¹Kyriba, Saint Cloud, France

²Université Pierre et Marie Curie, Paris, France

³Université de Reims Champagne Ardennes

January 2, 2017

Abstract

Opponent modeling consists in modeling the strategy or preferences of an agent thanks to the data it provides. In the context of automated negotiation and with machine learning, it can result in an advantage so overwhelming that it may restrain some casual agents to be part of the bargaining. We qualify as “curious” an agent driven by the desire of negotiating in order to collect information and improve its opponent model. However, neither curiosity-based rationality nor curiosity-robust protocol have been studied in automatic negotiation. In this paper, we rely on mechanism design to propose three extensions of the standard bargaining protocol that limit information leak. Those extensions are supported by an enhanced rationality model, that considers the exchanged information. Also, they are theoretically analyzed and experimentally evaluated.

Keywords: Automated Negotiation, Bargaining, Opponent Modeling, Rationality, Mechanism Design, Curiosity

1 Introduction

Negotiation involves several parties that exchange proposals in order to reach a mutually beneficial outcome. Each participant has private information and a utility function that may be in conflict with other participants.

To get a good agreement for herself, a participant may choose to use opponent modeling (see [3] for a survey on this topic). Opponent modeling is based on machine learning techniques, which are fed by the sequences of

*Corresponding Author: Cédric Buron, UPMC cedric.buron@lip6.fr
We thank Finance Innovation, ECM2, nov@log and the Région Île de France for supporting our project.

exchanged proposals. It allows to learn the way the opponent values goods ([6, 19]), which provides a major advantage in the future negotiations ([20]).

Opponent modelers have therefore an incentive to pretend to negotiate to get a good while they are seeking as much data as possible from their opponent. Opponent modelers that act this way are called “curious”. Curious agents have a harmful influence on negotiations: collected information provides them more efficient modelings, which will significantly unbalance the future negotiations in their favor. Consequently, casual agents may be discouraged and be reluctant to use the bargaining protocol.

In this paper, we design a protocol that extends the standard bargaining in order to restrain the sphere of operation of curious agents. To do that, we first present related work (??) and model how the agents evaluate exchanged information by extending the classical model of rationality (??). We then improve standard bargaining with three theoretically analyzed extensions (??). The first one prevents agents from initiating the negotiation if there is no chance to reach an agreement. The second one directly limits the number of exchanged proposals. The third one prohibits the agents from leaving the negotiation without agreement. The significance of our proposed extensions is proved empirically (??). We finally summarize our proposal and give some possible extensions to our work (??).

2 Related Work

In bargaining [4], each party tries to reach a joint decision involving two conflicting goals: maximizing the utility of an agreement and maximizing the chance of acceptance by the opponent. Relevant tactics have been proposed to achieve the first goal such as the efficient Tit-for-Tat [12]. The second goal requires skills such as curiosity to understand the opponents’ strategies and preferences. Psychologists Fisher *et al.* [9] define the curiosity as the ability to see the situation as the other side sees it.

In automated negotiation, this ability corresponds to what is called “opponent modeling” [3]. We distinguish two approaches: on single session (online) and on multi-session (offline). Among the work on single session modeling, the most widely used techniques are Bayesian learning [11], non-linear regression [13], and SVM [15]. Multi-session techniques include Kernel Density Estimation [19] and neural networks [5]. The preponderance of this approach is attested by the Automated Negotiation Agent Competition [1] since all high ranking agents make a wide use of it [2].

The benefits of opponent modeling can lead to manipulative behavior where agents try to get information at any cost. Also, the negotiation ends up at the expense of the agents that do not mean to invest in this technique. Bargaining being vulnerable to those behaviors, this protocol is disadvantageous for the latter and therefore unreliable in general case. It is thus

necessary to develop a new bargaining framework resistant to those practices. Such a topic belongs to the theory of mechanism design.

Mechanism design [18, 16] studies the conception of “incentive-compatible” mechanisms [14], *i.e.* protocols in which the agents act according to the desire of the designer of the protocol. As an example, in the context of mediated bargaining¹, several work (*e.g.* [17, 7]) give a theoretical framework on how to design a protocol in which the agents provide their real reserve prices to the mediator.

To the best of our knowledge, there is no work on mechanism design related to standard (non-mediated) bargaining. In particular, there is no attempt to design an extension of bargaining resistant to the aforementioned problem.

3 Model of Curiosity-Aware Bargaining

We introduce in this section an original rationality model for bargaining. Agents consider not only the price associated to the good but also the leaked and collected information. We first introduce a model of the bargaining protocol that explicitly represents this information. Then we extend the two affected objects: the agent utility and the reserve price. Finally we propose a classification of curiosity-aware agents.

3.1 Bargaining Model

Bargaining is a type of negotiation in which two agents, a *purchaser* p and a *seller* s , seek mutual agreement on the price of a good. Each agent is endowed with an **agreement domain** that induces, for the seller, the minimum price it is willing to sell the good, and for the purchaser, the maximum price it is willing to pay for it. Those values are called *reserve prices*; an agent has a null utility if the agreed price is exactly the reserve price, a positive utility if it is better and a negative one if it is worse.

A bargaining proceeds with a **a sequence of exchanged proposals** which respect the monotonic concession principle: in any two successive proposals of an agent the second is always better for it opponent than the first. A bargaining **ends either by an accept or a reject** by one of the agents. The reject may be sent at any time to end the negotiation without any transaction.

Definition 1 (Bargaining). *A bargaining proceeds with exchange of proposals and ends with an accept or a reject of one of the participants. We note V the space of lists of messages sent by one party during a bargaining.*

$$V = \{(k, p, e) | k \in \mathbb{N}, p \in \mathbb{R}^k, e \in \{\emptyset, \top, \perp\}\} \quad (1)$$

¹in mediated bargaining, the negotiators use a trusted third party called mediator. The latter uses information delivered by the negotiators to generate the proposals.

where

$k \in \mathbb{N}$ is the number of proposals

$p \in \mathbb{R}^k$ is the list of proposed prices

$e \in \{\emptyset, \top, \perp\}$ is the ending message: either \top if this party accepts the last deal, \perp if it rejects it or \emptyset if the other party ends the negotiation.

We note \mathcal{B} the set of bargaining records. A bargaining $b \in \mathcal{B}$ is defined with a six-tuple:

$$b = (g_b, s_b, p_b, \pi_b, \nu_b^s, \nu_b^p) \quad (2)$$

where

$g_b \in \mathcal{Good}$ is the negotiated good,

$s_b, p_b \in \mathcal{Agent}^2$ are respectively the seller and the purchaser identifiers

$\pi_b \in \mathbb{R}^+ \cup \{\perp\}$ is the negotiated price if the negotiation succeeds or \perp if it fails.

$\nu_b^s, \nu_b^p \in V^2$ are respectively the sequence of the seller and purchaser proposals.

3.2 Curiosity-Aware Utility & Reserve Price

In the classical rationality model used in bargaining, an agent is only interested in the final price of the good. A curiosity-aware agent is also interested in the information exchanged during the bargaining process. This affects both its utility model and its reserve price.

Definition 2 (Curiosity-Aware Utility). *The utility of a curiosity-aware agent is computed from the final price and the sequence of exchanges.*

$$\forall a \in \mathcal{Agent}, \forall b = (g_b, s_b, p_b, \pi_b, \nu_b^s, \nu_b^p) \in \mathcal{B}, \quad u_a : \mathbb{R} \times V \times V \rightarrow \mathbb{R} \\ (\pi_b, \nu_b^s, \nu_b^p) \mapsto u_a(\pi_b, \nu_b^s, \nu_b^p) \quad (3)$$

The reserve price is the price for which the agent utility is null. Classical agents have a constant reserve price throughout the negotiation, while the reserve price of curiosity-aware agents is affected by the proposals made so far.

Definition 3 (Curiosity-Aware Reserve Price). *Let b be a bargaining. We note $\pi_b^p(k) \in \mathbb{R}$ the reserve price of the purchaser at the intermediary step k . It is defined by the following property:*

$$\forall b \in \mathcal{B}, \forall k \leq k_b^p, u_p(\pi_b^p(k), \nu_{\setminus k}^s, \nu_{\setminus k}^p) = 0 \quad (4)$$

where k_b^p is the number of proposals made by purchaser p (see ??), $\nu_{\setminus k}^p$ holds the k first proposals of the list ν^p and $\nu_{\setminus k}^s$ holds the k first proposals of the list ν^s (if the seller made the first proposal, $k - 1$ else).

3.3 Agent Taxonomy

We now introduce three canonical types of agents that are characterized by the way they value information and price. For the sake of simplicity, we rely on comparators that are indifferently used for sellers and purchasers. For Agent a :

\succ_a^π **and** \sim_a^π are respectively strict and equivalence orders over prices. A seller prefers a higher price whereas a purchaser prefers a lower one.

\succ_a^ν **and** \sim_a^ν are respectively strict and equivalence orders over sequences of proposals (*i.e.* V). For two sequences of proposals ν and ν' , $\nu \succ_a^\nu \nu'$ means that agent a considers that ν carries more information than ν'

In order to consistently extend the classical bargaining rationality model, we assume that, any other thing being equal (*ceteris paribus*), any agent prefers a bargaining resulting in a better price outcome. In other words, if two bargainings reveal the same amount of information, an agent a prefers the one that optimizes the price. The following equation is verified by any rational agent:

$$\forall b, b' \in \mathcal{B}^2, \left\{ \begin{array}{l} \pi_b \succ_a^\pi \pi_{b'} \\ \nu_b^s \sim_a^\nu \nu_{b'}^s \\ \nu_b^p \sim_a^\nu \nu_{b'}^p \end{array} \right\} \implies u_a(\pi_b, \nu_b^s, \nu_b^p) > u_a(\pi_{b'}, \nu_{b'}^s, \nu_{b'}^p) \quad (5)$$

Curiosity-aware agents may also consider both the leaked and revealed information. An agent can be *secretive* or *unsecretive* depending on whether or not it is concerned about the information it reveals. Also, an agent can be *curious* or *uncurious* depending on whether or not it is interested in the information it collects. ?? presents examples of reserve price functions for the three types of agents we detail below. They are drawn from ??, used in our experiments.

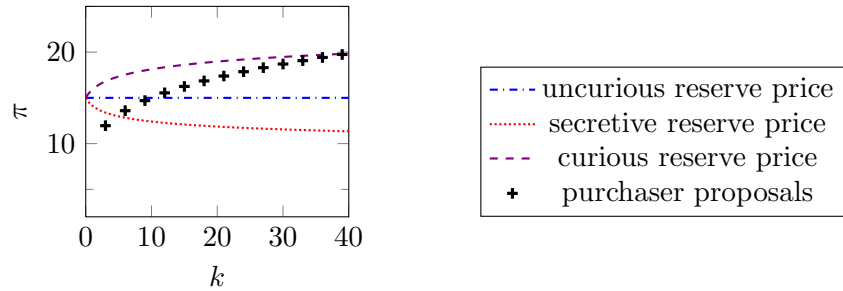


Figure 1: Proposals and reserve price of uncurious, curious and secretive agents

An Uncurious (unsecretive) agent implements the classical understanding of rationality in bargaining: its sole interest is the final price.

Definition 4 (Uncurious Agent). *An uncurious agent utility is not affected by the exchanged information.*

$$\forall b, b' \in \mathcal{B}^2, \pi_b \succ_a^\pi \pi_{b'} \implies u_a(\pi_b, \nu_b^s, \nu_b^p) > u_a(\pi_{b'}, \nu_{b'}^s, \nu_{b'}^p) \quad (6)$$

A Secretive (uncurious) agent is aware that its actions reveal information which can be used by malicious agents at its expense. It is however only interested in selling/purchasing goods and not in collecting information about other agents. The price it is willing to pay for a good decreases while it is making proposals. Thus, it tends to drop the negotiation sooner than a classical (uncurious) agent. Also, if the deal is not concluded, its utility is negative since it reveals information and does not get any outcome.

Definition 5 (Secretive Agent). *A secretive agent prefers not to reveal information, but is not interested by the collected information. In the case of a purchaser p :*

$$\forall b, b' \in \mathcal{B}^2, \left\{ \begin{array}{l} \pi_b \sim_p^\pi \pi_{b'} \\ \nu_b^p \succ_p^\nu \nu_{b'}^p \end{array} \right\} \implies u_p(\pi_b, \nu_b^s, \nu_b^p) > u_p(\pi_{b'}, \nu_{b'}^s, \nu_{b'}^p) \quad (7)$$

A Curious (unsecretive) agent is interested in collecting information that it may use for opponent modeling. It is not concerned about the information it may reveal. The price it is willing to pay for a good increases when the amount of collected information increases. It has a positive utility in case of reject. A good strategy for such an agent is to make the bargaining end as late as possible and to eventually reject it.

Definition 6 (Curious Agent). *A curious agent aims collecting as much information as possible during the exchange, but is not concerned by leaked information. In the case of a purchaser p :*

$$\forall b, b' \in \mathcal{B}^2, \left\{ \begin{array}{l} \pi_b \sim_p^\pi \pi_{b'} \\ \nu_b^s \succ_p^\nu \nu_{b'}^s \end{array} \right\} \implies u_p(\pi_b, \nu_b^s, \nu_b^p) > u_p(\pi_{b'}, \nu_{b'}^s, \nu_{b'}^p) \quad (8)$$

A Curious and Secretive agent is both doing opponent modeling and aiming to preserve itself from being modeled by other agents. It implements ?? which is a more constrained version of ???? (noted (??') and (??')). It is interested by either collecting more information *ceteris paribus*, or revealing less information *ceteris paribus*. In the case of a purchaser p :

$$\forall b, b' \in \mathcal{B}^2, \left\{ \begin{array}{l} \overset{(\text{??'})}{\pi_b \sim_p^\pi \pi_{b'}} \\ \overset{(\text{??'})}{\nu_b^s \sim_p^\nu \nu_{b'}^s} \\ \nu_b^p \succ_p^\nu \nu_{b'}^p \end{array} \right\} \vee \left\{ \begin{array}{l} \overset{(\text{??'})}{\pi_b \sim_p^\pi \pi_{b'}} \\ \overset{(\text{??'})}{\nu_b^p \sim_a^\nu \nu_{b'}^p} \\ \nu_b^s \succ_a^\nu \nu_{b'}^s \end{array} \right\} \implies u_p(\pi_b, \nu_b^s, \nu_b^p) > u_p(\pi_{b'}, \nu_{b'}^s, \nu_{b'}^p) \quad (9)$$

We highlight the fact that the three previously introduced agent types can be used as atomic building blocks of this last type. Indeed, during an exchange of messages, at each step, a curious and secretive agent acts like one and only one canonical type. For instance, when comparing two negotiations where $(??')$ (resp. $(??')$) applies, its utility is similar to the one of a secretive (resp. curious) agent. In the other numerous cases where $??$ is not applicable, it can be abstracted to a secretive agent each time it considers that the leaked information is more critical than the collected one, as a curious each time it values more the collected information than the revealed one, or as an uncurious if it estimates than the revealed and collected information compensate each other.

4 Protocols for Curiosity-Proof Bargaining

To limit the impact of curious agents, we introduce and analyze variants of the standard bargaining protocol.

We first analyze the different flaws of classical bargaining *w.r.t.* curious agents and then introduce a curiosity-proof extension for each phase of the protocol. Last, we theoretically analyze some induced incentives.

As described in $??$, bargaining can be analyzed in term of three aspects: (1) the agreement domains, (2) the sequence of proposals and (3) the way the negotiation ends (accept or reject). Each of those aspects may lead to situations with unnecessary leak of information, which may be exploited by curious agents.

The mutual agreement domain is the range of prices where a mutual agreement can be found. If it is empty (the seller minimal price is higher than the purchaser maximal price) the bargaining is not accepted: the revealing of information is useless. Furthermore, a curious seller may decide of a reserve price much higher than its real one. If the bargaining is concluded, it would collect information and sell the good at a very high price, which is very good for it; otherwise it would collect information for free.

The sequence of proposals is defined by the agent strategies and reserve prices. Adopting strategies that make smartly chosen concessions can increase the number of exchanges. Such strategies can be used by a curious agent in order to optimize the collected information.

The negotiation issue (accept or reject) is also a crucial aspect. A curious agent gets a positive utility as soon as some proposals are done, even if the negotiation ends with a reject. *A contrario*, in the case of a reject, the utility of a secretive agent is negative. This puts the ratio of power strongly in favor of the curious agents who may accept only if the proposed price is really advantageous.

We now present a variant for each of the steps that aims to reduce curiosity.

4.1 Description of the proposed modifications

Matching We first propose to avoid a useless leak of information resulting from bargaining between two parties that have incompatible agreement domains. A trusted third-party collects the agents (initial) reserve prices and authorizes the bargaining if and only if a mutual agreement is possible.

Bounds Secondly, we propose to bound the number of exchanges ; if no agreement has been reached by the bound, the bargaining fails. The bound associated to a bargaining $b \in \mathcal{B}$ is noted k_b^{max} . Gneezy *et al.* [10] shows that this advantages the agent making the last proposal ; similarly to [21] we rely on a trusted third-party which transmits the proposals simultaneously, after both agents have made them.

Enforcing the execution of a transaction Lastly, we propose to prevent the possible failure of a bargaining by enforcing the transaction even if the deal has been rejected or the bound of proposals has been reached. In this case, the good is exchanged at the agents reserve prices and the remaining money is taken by the community as a penalty. This transaction is the less favorable for both parties and thus provides a strong incentive to conclude all the deals.

4.2 Some proven properties

We now prove two properties about the incentives induced by our proposals. First, ?? shows that the proposals provide a strong incentive, for a purchaser (resp. seller), not to declare a too high (resp. too low) reserve price. This means that, for instance, a curious agent will not lie on its reserve price in order to match as many agents as possible and cancel the negotiation after it has got enough information. Note that the question of providing an incentive, for a purchaser (resp. seller) not to declare a reserve price lower (resp. higher) than the real one is still open. Hence, it is possible to manipulate the matching procedure in order to detect adversary parties with advantageous reserve prices and with the guarantee of getting a profitable deal.

Theorem 1. *Suppose the agent considers that the amount of information increases with the number of exchanged messages. Then, using the three extensions gives a curious purchaser (resp. seller) an incentive to give a reserve price inferior (resp. superior) to $\pi_b^p(k_b^{max})$ (resp. $\pi_b^s(k_b^{max})$).*

Proof. Without loss of generality, we suppose that our agent is a curious purchaser p_b . Let b be a bargaining (see ??) between p_b and a seller s_b . We prove that if the purchaser declares a more interesting price (for the seller) than its reserve price computed at the bound, either it pays more (case 1 and 2) or it is matched with the same seller as if it was honest (case 3).

Case 1: $\pi_b = \perp$, the deal is rejected. In this case, due to the third extension, the agent pays the price it has declared. Hence, the lower the price, the better the deal. In particular, if the bound is reached, p_b has to pay price higher than its reserve price and gets a negative utility.

Case 2: $\pi_b \in \mathbb{R} \wedge \pi_b > \pi_b^p(k_b^{max})$, the bargaining is accepted, but p_b pays a price higher than his true reserve price. This case may arise if the last proposal before the bound is higher than its true reserve price, but lower than the declared one. It is then more interesting for p_b to accept the proposal than to reach the bound and pay the declared reserve price. Such a situation can happen whether the seller has a reserve price higher or lower than $\pi_b^p(k_b^{max})$. In the first case, p_b is matched with an agent and would have not be matched with it if it was honest, but it pays more than it wants; in the second it would have been matched with this agent anyway.

Case 3: $\pi_b \in \mathbb{R} \wedge \pi_b^p(k_b^{max}) \geq \pi_b$, the bargaining is accepted at a price lower than the true reserve price of p_b . This is the only case where p_b gets a good at an acceptable price. However, assuming the seller honesty, this can only occurs if the true reserve price of p_b is compatible with the reserve price of s_b . Hence, s_b would have been matched with p_b even if p_b had been honest.

□

?? shows that the proposals also provide a strong incentive to find an agreement while negotiating.

Theorem 2. *Rejecting an offer or going to a deadline is less interesting for a couple of rational negotiating agents than finding any agreement more advantageous than the declared reserve price, before the negotiation ends.*

Proof. Suppose that the bargaining b failed. The utility of the purchaser is $u_p = (\pi_p^{matching}, k_b^s, k_b^p)$ where $\pi_p^{matching}$ is the reserve price that the purchaser has declared. According to ??, p would prefer any situation with a success on a price inferior than $\pi_p^{matching}$. The same goes for the seller. □

5 Effect of the introduced extensions on the curiosity

In this section we propose an empirical analysis of the impact of the introduction of the proposed extensions. Our objective is to measure how an inappropriate collect of information is restrained. Restraining the collect of information without impacting the protocol performance ideally reduces the welfare of curious agents, increases the one of secretive agents and does not affect uncurious agents.

We first describe the experimental protocol. Then, we present empirical results about the variation of the number of exchanged messages and the agent welfare.

5.1 Experimental Protocol

Our experiments consist in running different instances of bargaining with randomly generated couples of agents. A bargaining instance is identified to the possible variants: standard bargaining, with matching, with a given bound or with the three extensions altogether. The bound on the number of messages, k_b^{max} , is set to 500. This value has been computed from a set of experiments not presented here for reasons of space. It is the most advantageous for secretive agents on our specific experimental context: below this value, we observed that the bound strongly decreased their welfare, whereas above this value it remained constant. The same phenomenon was observed for uncurious and curious agents with, respectively, the values 750 and 1000. Those values correspond to the point where most of the deals were already ended.

An agent a is associated to a type (*uncurious*, *curious* or *secretive*), an initial reserve price $\pi_a(0)$, a utility function u_a , and a strategy function. The utility function of the agent is represented² in ???. One can verify the validity of those functions (see ???) *w.r.t.* ?????? .

$$\forall b \in \mathcal{B}, u_a(\pi_b, \nu_b^s, \nu_b^p) = \begin{cases} \pi_a^b(0) \cdot \Delta(n) - \pi_b & \text{if the negotiation succeeds} \\ \pi_a^b(0) \cdot (\Delta(n) - 1) & \text{else} \end{cases}$$

where

$$\begin{cases} \Delta(n) = 1 & \text{if } a \text{ is uncurious} \\ \Delta(n) = \frac{1}{\log_n(n + k_b^p)} & \text{if } a \text{ is secretive} \\ \Delta(n) = \log_n(n + k_b^p) & \text{if } a \text{ is curious} \end{cases} \quad (10)$$

A strategy function indicates at each step the proposal the agent is willing to make. If the other party makes a better proposal, the agent accepts

²Only the utility of purchaser is given here, the utility of seller being similar

it. If the planned proposal is worse than the agent reserve price, the agent drops off the bargaining. The used strategy is represented in ?? and is an adaptation of [8]. The parameters κ_a , β_a and γ_a determine the shape of the curve for agent a . An example is plotted in ??.

$$\gamma + \left(\kappa_a + (1 - \kappa_a) \cdot \left(\frac{k}{k_a^{max}} \right)^{\frac{1}{\beta_a}} \right) \cdot (\gamma_a - \pi_a^b) \quad (11)$$

An experiment considers a couple of agents, defined by their type and an instance of bargaining. The agent reserve prices have been drawn from a Gaussian distribution and κ_a , β_a and γ_a from uniform distributions. Each experiment is averaged over 10'000 draws and we observe the final welfare of each type of agent.

5.2 Impact on the welfare

?? displays the welfare of the different types of agents in different settings. “barg” stands for standard bargaining, “mat.” for the matching extension, “bou.” for the bounding extension (with a bound of 500 messages) and “all” for the three extensions all together.

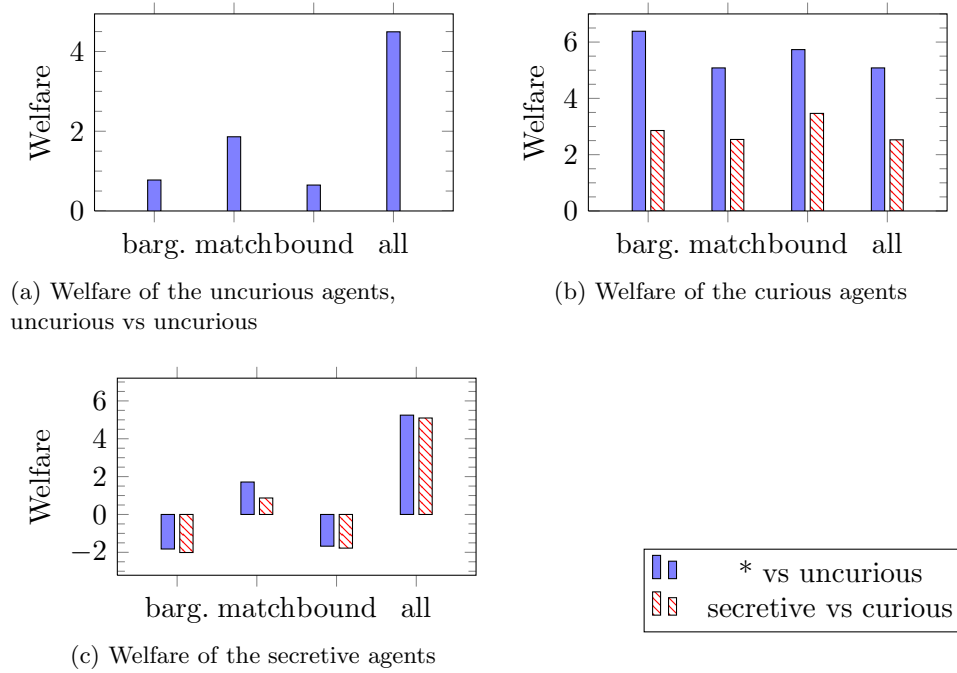


Figure 2: Welfare of the agents in different situations

?? shows how the proposed extensions affect the welfare of classical agents. Bounding slightly reduces their welfare (10%): it tends to close the

longest negotiations by rejects. *A contrario*, all others increase noticeably their welfare. Matching doubles it by increasing the number of successful negotiations. Applying all extensions quadruples it: the agents make concessions more quickly in order to reach their reserve price before the deadline (??).

?? shows the impact on curious agents. First, we observe that their welfare is worse when they negotiate with secretive agents rather than uncurious ones. Indeed, secretive agents tend to drop off the bargaining more often. Whatever the opponent, matching and applying all extensions lower the welfare by a sixth since agents get to an agreement more quickly. *A contrario*, using only bounding lowers the welfare by 10% when negotiating with uncurious agents while it increases it by the same factor when negotiating with secretive ones. Please note that (1) bounding alone results in more reject, thus curious agents get more often information “for free” and (2) curious agents might accept overpriced deal for the sake of information collect. This last point is more critical when facing secretive agents since they get more tough as the bargaining is going on. Consequently, if they were to succeed, the deals rejected thanks to the bound would have been more disadvantageous for curious agents when they are facing secretive agents than with uncurious ones.

?? shows the impact of the extensions on secretive agents. Each extension has a positive impact. Please note that the average welfare of secretive agents is negative with the standard protocol because of the negotiations that end by a reject: secretive agents have revealed information without counterpart. Matching limits considerably the number of rejects, giving thus a positive welfare to them with nearly the same absolute value. Bounding only slightly increases the welfare of secretive agents (its absolute value is reduced by 9%) by reducing the number of exchanged messages before a reject. Last, adding the enforcement of transaction results in a really higher welfare (three times in absolute value) for the same reason as uncurious agents. The welfare of secretive agents is always worse when they negotiate with curious agents.

The extensions that we have introduced reach their goal by preventing curiosity. Both matching agents on their reserve prices and forcing them to make deals effectively make opponent modeling much harder. However, bounding the number of exchanged messages alone is not very beneficial.

6 Conclusion and Perspectives

This article studied curiosity in bargaining. Curious agents exploit information leaked during a bargaining in order to model their opponents and, thus, gain a significant advantage in future bargaining. In order to reduce

it, we first introduced an original rationality model that considers the exchanged information. This led to the establishment of three types of agents: uncurious agents that correspond to the classical agents – they do not consider the exchanged information; secretive agents that are concerned by the information they might unveil; and curious agents that seek for information leaked by their opponents. We then proposed three extensions of the standard bargaining protocol both to limit the exchanged information and to provide incentives to not seek it. Incentives have been studied theoretically and the general impact of those extensions has been analyzed empirically.

Our experiments showed how strongly curiosity impacts the bargaining process. Indeed, introduction of curious agents results in longer negotiation and higher prices, while it is the contrary for secretive ones. Used simultaneously, our three extensions have proven to be really effective for preventing agents from seeking information, assuming they declare their true intention. However, this extended bargaining protocol can be manipulated by agents that are willing to negotiate less often but in more favorable conditions.

We plan to enrich our proposal in future work by addressing this last issue in the context of one-to-many negotiation. Indeed, the competitive environment of, for instance, Iterated Contract Net Protocol would help to deal with this issue.

References

- [1] The automated negociation agent competition website.
http://mmi.tudelft.nl/negotiation/index.php/Automated_Negotiating_Agents_Competition
- [2] Tim Baarslag, Reyhan Aydoğan, Koen V Hindriks, Katsuhide Fujita, Takayuki Ito, and Catholijn M Jonker. The automated negotiating agents competition, 2010–2015. *AI Magazine*, 34(6):115–118, 2015.
- [3] Tim Baarslag, Mark J C Hendrikx, Koen V Hindriks, and Catholijn M Jonker. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems*, pages 1–50, 2015.
- [4] Lawrence E Blume and Steven F Durlauf, editors. *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, 2008.
- [5] Réal Carbonneau, Gregory E. Kersten, and Rustam Vahidov. Predicting opponent’s moves in electronic negotiations using neural networks. *Expert Systems with Applications*, 34(2):1266–1273, 2008.
- [6] David Carmel and Shaul Markovitch. Opponent modeling in multi-agent systems. In *Adaption and learning in multi-agent systems*, pages 40–52. Springer, 1996.

- [7] Kaylan Chatterjee. Incentive compatibility in bargaining under uncertainty. *The Quarterly Journal of Economics*, 97(4):717–726, 1982.
- [8] Peyman Faratin, Nicholas R Jennings, and Carles Sierra. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 24(3-4):159–182, 1998.
- [9] Roger Fisher, William L Ury, and Bruce Patton, editors. *Getting to yes: Negotiating agreement without giving in*. 2d ed. New York: Penguin, 1991.
- [10] Uri Gneezy, Ernan Haruvy, and Alvin E Roth. Bargaining under a deadline: evidence from the reverse ultimatum game. *Games and Economic Behavior*, 45(2):347–368, 2003.
- [11] Koen Hindriks and Dmytro Tykhonov. Opponent modelling in automated multi-issue negotiation using bayesian learning. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, volume 1, pages 331–338. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
- [12] Koen V Hindriks, Catholijn M Jonker, and Dmytro Tykhonov. The benefits of opponent models in negotiation. In *IEEE/WIC/ACM International Joint Conferences on Web Intelligence and Intelligent Agent Technologies, 2009. WI-IAT '09.*, volume 2, pages 439–444, Sept 2009.
- [13] Chongming Hou. Modelling agents behaviour in automated negotiation. Technical report, Knowledge Media Institute, 2004.
- [14] Leonid Hurwicz. *Decision and organization*, chapter On informationally decentralized systems, pages 297–336. North-Holland Pub. Co., 1972.
- [15] Kennard Lavers, Gita Sukthankar, Matthew Klenk, David W Aha, and Matthew Molineaux. Opponent modeling and spatial similarity to retrieve and reuse superior plays. In *Proceedings of the ICCBR Workshop on Case-Based Reasoning for Computer Games*, 2009.
- [16] Eric S Maskin. Mechanism design: How to implement social goals. *The American Economic Review*, 98(3):567–576, 2008.
- [17] Roger B Myerson. Incentive compatibility and the bargaining problem. *Econometrica*, 47(1):61–73, 1979.
- [18] Roger B Myerson. *Allocation, Information and Markets*, chapter Mechanism design, pages 191–206. Palgrave Macmillan UK, 1989.

- [19] Yinon Oshrat, Raz Lin, and Sarit Kraus. Facing the challenge of human-agent negotiations via effective general opponent modeling. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 377–384. International Foundation for Autonomous Agents and Multiagent Systems, 2009.
- [20] David D B Van Bragt and Han La Poutr . Why agents for automated negotiations should be adaptive. *Netnomics*, 5(2):101–118, 2003.
- [21] LinLan Zhang and Qing Liu. An automated multi-issue negotiation mechanism based on intelligent agents in e-commerce. *Journal of Advanced Management Science*, 4(2):172–175, Mar 2016.